

SWITCHING DEVICE INTERFACES

Background of the Invention

This invention relates generally to interfaces on a network device.

5 Network devices have physical interfaces that are subject to failure. When such an interface fails, a network device can be cut-off from the network. This is particularly problematic in the case of a router, where failure of a single physical interface can make a whole
10 branch of the network inaccessible to other devices.

Summary of the Invention

15 In one aspect of the invention, a device is switched from a first physical interface on the device (for example, a failed interface) to a second physical interface on the device based on information in an interface redundancy group. The information in the interface redundancy group identifies the first physical interface as a primary interface for the device and the second physical interface as a secondary interface for the device.

20 The foregoing aspect of the invention may include one or more of the following features/functions.

 The interface redundancy group may include information defining the primary interface for the device and one or more secondary interfaces for the device. An
25 event may be detected at the first physical interface, and switching may be performed in response to the event. The event may comprise a failure of the first physical interface. The first physical interface may be associated

with a driver and a signaling stack, and the failure of the first physical interface may comprise a failure of the driver and/or the signaling stack. The driver and the signaling stack may be monitored in order to detect failures
5 therein. The event may comprise receipt of a slot failure at the first physical interface.

Prior to switching, the second physical interface may operate in a passive mode during which the second physical interface is dormant. Prior to switching, the
10 second physical interface may operate in an active mode during which the second physical interface is communicating over a network. The first physical interface may support one or more network layer interfaces. Following switching, the second physical interface may support the one or more
15 network layer interfaces formerly supported by the first physical interface. The first and second physical interfaces may comprise asynchronous transfer mode ("ATM") physical interfaces. The first and second physical interfaces may be resident on a single network router.

Following switching, the second physical interface may assume responsibilities of the first physical interface. These responsibilities may include routing and/or bridging functions. Following switching, the second physical interface may be configured in a same manner as the first
20 physical interface was configured prior to switching. The device may include a third physical interface, and the interface redundancy group may identify the third physical interface as a tertiary interface. The device may be switched from the second physical interface to the third
25 physical interface in response to an event. Following switching, the third physical interface may assume responsibilities of the first and second physical
30

interfaces. These responsibilities may include routing and/or bridging functions.

In another aspect, physical interfaces on a single device are switched by designating a physical interface on the device as a high priority physical interface, and determining if the high priority physical interface is available. The device is switched from a lower priority physical interface to the high priority physical interface when the high priority physical interface is available.

The foregoing aspect of the invention may include one or more of the following features/functions. Switching may be performed automatically in response to the high priority interface being available. The high priority physical interface may be monitored to determine if the high priority physical interface is available.

In another aspect, a device is switched from a first ATM physical interface on the device to a second ATM physical interface on the device based on information in an interface redundancy group. The information in the interface redundancy group identifies the first ATM physical interface as a primary interface for the device and the second ATM physical interface as a secondary interface for the device. ATM network layer interfaces are established over the second physical interface that correspond to ATM network layer interfaces that were established over the first ATM physical interface prior to switching.

This brief summary has been provided so that the nature of the invention can be understood quickly. A detailed description of illustrative embodiments of the invention is set forth below.

Brief Description of the Drawings

FIG. 1 shows a network system, which includes a router having switchable physical interfaces.

FIG. 2 shows virtual circuits supported by the physical interfaces on the router.

FIG. 3 is a flow diagram showing a process for switching a physical interface on the router in "passive" mode.

FIG. 4 shows an alternative connection of the virtual circuits to the physical interfaces on the router.

FIG. 5 is a flow diagram showing a process for switching a physical interface on the router in "active" mode.

FIG. 6 shows the configuration of the virtual circuits after the switching performed in FIG. 5.

FIG. 7 is a flow diagram showing a process for switching a physical interface on the router based on priority.

FIG. 8 shows a router having three physical interfaces and redundancy groups therefor.

Description of the Preferred Embodiment

Referring to FIG. 1, a network system 10 is shown. Network system 10 includes router 12, switches 14 and 16, workstations 18, 20 and 22, and network 24.

Network 24 is an asynchronous transfer mode ("ATM") wide area network ("WAN"). ATM is a connection-oriented protocol, meaning that connections are established between devices before data and/or communications can be transmitted between the devices. The network layer interface comprises virtual circuits, over which data packets (in ATM parlance, "cells") are transmitted among devices coupled to network 24, such as router 12 and switches 14 and 16. Virtual

circuits can be established by protocols such as ELANs (Emulated Local Area Networks) on network 24.

Switches 14 and 16 interface workstations 20 and 22, respectively, to network 24. Each switch 14 and 16 is an electronic device that routes cells between network 24 and a corresponding workstation. Workstations 18, 20 and 22 are personal computers ("PCs") or other devices that are capable of receiving cells from network 24, processing cells, and transmitting cells to network 24.

Router 12 is a computer or other device that transmits packets/cells among workstations 18, 20 and 22 via network 24. For example, router 12 receives cells/packets from workstation 18 and, based on information in those cells/packets, routes the cells/packets to either workstation 20 or 22 (through an intervening switch and other hardware on network 24).

ATM physical interfaces 26a and 26b are ports provided on router 12 for interfacing to network 24. Although only two such interfaces are shown, any number may be provided. In system 10 of FIG. 1, one of the interfaces (e.g., 26a) is designated as the primary interface for router 12 and the other (e.g., 26b) is designated as the secondary (or backup) interface. When the primary physical interface 26a fails, the secondary physical interface 26b is switched in to take its place. A process to accomplish this is described below.

Physical interfaces 26a and 26b are comprised of wires terminating in connectors, such as an OC-3/OC-12 connector, which mate to corresponding receptacles on router 12. One or more network layer interfaces are established by router 12 over each physical interface 26a and 26b for communication to, and over, network 24. ATM line drivers

30a and 30b transmit cells over corresponding physical interfaces 26a and 26b.

Included in router 12 are a processor 34 and a memory 36 connected by bus 38 (see view 40). Memory 36 stores routing engines 42a and 42b, interface switching code 44, and signaling stacks 46a and 46b. Processor 34 executes instructions in this code to cause router 12 to perform the functions described below. Memory 36 also stores interface redundancy group information 48 (described below).

Signaling stacks 46a and 46b are blocks of code, associated with corresponding physical interfaces 26a and 26b, for establishing virtual connections for the network layer interfaces over the physical interfaces. As noted, one or more network layer interfaces may be configured over a single ATM physical interface 26a and 26b.

Routing engine 42a routes cells over physical interface 26a and routing engine 42b routes cells over physical interface 26b. Routing engines 42a and 42b examine destination information in the cells and route the cells over appropriate interfaces. Examples of routing engines that may be used are "ARE" (ATM Routing Engine) and the 5782 Centillion Multiprotocol Engine.

Interface redundancy group information 48 defines which physical interface is the primary interface (e.g., 26a) and which is the secondary interface (e.g., 26b). This information may be input manually at router 12 via configuration software such as Site Manager® or Bay Command Console® ("BCC"). This software is used by network administrators to configure network devices. Alternatively, interface redundancy group information 48 may be downloaded from a remote location such as network 24 or workstation 18 or set via interface switching code 44.

Interface redundancy group information 48 includes user-defined redundancy groups. These redundancy groups assign priority to the interfaces. In the case of a two-interface router, such as router 12, there are two possible groups. For example, physical interface 26a is configured as the primary interface and physical interface 26b is configured as the secondary interface. The routing engine for each interface is configured to know the role of the interface in each redundancy group. Representative code to configure routing engines is provided in the Appendix.

In routers with more than two physical interfaces, interface redundancy groups become more complicated. Table 1 shows an example of redundancy groups for a router having four physical interfaces "Interface 1", "Interface 2", "Interface 3" and "Interface 4" (not shown).

Table 1

Redundancy Group	Primary Interface	Secondary Interface	Tertiary Interface
1	Interface 1	Interface 2	Interface 3
2	Interface 2	Interface 1	Interface 3
3	Interface 4	Interface 3	--

For example, in redundancy group "1", "Interface 1" acts as the primary interface; "Interface 2" acts as the secondary (first backup) interface and is used if "Interface 1" fails; and "Interface 3" acts as the tertiary (or second backup) interface and is used if both "Interfaces 1" and "Interface 2" fail.

Detecting physical interface failure and switching from a primary to a secondary physical interface (or from a

secondary to a tertiary physical interface, etc.) is performed by interface switching code 44. The operation of interface switching code 44 differs depending upon whether the secondary interface is in passive mode or active mode.

5 Passive Mode

In passive mode, prior to switching, the secondary interface is dormant. That is, the secondary interface is not driving/receiving signals to/from network 24. Passive mode may be set as the default mode of router 12 using Site Manager® or BCC®.

Referring to FIG. 2, a graphical representation of passive mode is shown. In FIG. 2, physical interface 26a is configured as the primary interface and physical interface 26b is configured as the secondary interface in passive mode. This configuration is set in interface redundancy group information 48. In FIG. 2, lines 54a to 54h represent connections maintained by primary physical interface 26a and lines 56a to 56h represent connections that are maintained by secondary physical interface 26b after secondary interface 26b takes over the role of primary interface 26b.

Referring to FIG. 3, an interface switching process 56 is shown. Interface switching process 56 is performed by interface switching code 44 to switch from primary physical interface 26a to secondary physical interface 26b. Process 56 monitors 58 primary physical interface 26a, including both driver 30a and signaling stack 46a, for specific "events". These events can include, but are not limited to, receipt of a slot reset at primary physical interface 26a and/or a failure of primary physical interface 26a, such as a failure of driver 30a and/or signaling stack 46a.

In response to detecting 60 one of the foregoing events, process 56 switches 62 from primary physical interface 26a to secondary physical interface 26b (in accordance with interface redundancy group information 48).

5 Switching 62 is performed by enabling driver 30b for secondary physical interface 26b. Following switching, secondary interface 26b establishes 63 the network layer interfaces to network 24 (in this example, ELANs 52a to 52h), and assumes the responsibilities (including routing and bridging services) of primary interface 26b. Switching
10 between interfaces is performed as quickly as possible, e.g., within thirty seconds of failure or reset.

With driver 30b enabled, in 64, signaling stack 46b transmits cells over secondary physical interface 26b to
15 ELANs 52a to 52h. Primary physical interface 26a can be repaired while secondary physical interface 26b performs its functions.

Active Mode

In active mode, prior to switching, the secondary
20 physical interface 26b is already communicating over the network. Following failure of the primary physical interface 26a, the secondary interface 26b assumes the responsibilities of the primary interface 26a. In particular, upon switching, code in the signaling stack
25 establishes the network layer interfaces (e.g., over ELANs) of the primary interface 26a, and assumes the routing and bridging functions of the primary interface 26a.

Referring to FIG. 4, a graphical representation of active mode is shown. In FIG. 4, physical interface 26a is
30 the primary interface and physical interface 26b is the secondary interface in active mode. This configuration is set in interface redundancy group information 48. Since

secondary physical interface 26b is in active mode, prior to switching, it is providing network layer services over virtual circuits for ELANs 52e to 52h. Primary physical interface 26a is supporting ELANs 52a to 52d.

5 Referring to FIG. 5, an interface switching process 66 is shown, that is performed by interface switching code 44 to switch from primary physical interface 26a to secondary physical interface 26b. Interface switching process 66 monitors 68 primary physical interface 26a,
10 including both driver 30a and signaling stack 46a, for specific events. These events are the same as those noted above for passive mode.

In response to detecting 70 one of the foregoing events, process 66 switches 72 from primary physical
15 interface 26a to secondary physical interface 26b. Switching is performed in the manner described above; that is, enabling and disabling the drivers for the appropriate network interfaces.

Interface switching process 66 establishes 74, over
20 secondary physical interface 26b, the network layer interfaces (in this case, ELANs 52a to 52d) of primary physical interface 26a. Secondary physical interface 26b continues to perform its original routing and bridging functions over ELANs 52e to 52h. Now, however, secondary
25 physical interface 26b also performs the routing and bridging functions formerly performed by primary physical interface 26a over ELANs 52a to 52d.

With driver 30b enabled, and secondary physical interface 26b switched, in 76, signaling stack 46b transmits
30 cells over secondary physical interface 26b to all of ELANs 52a to 52h. This is shown in FIG. 6. Primary physical interface 26a no longer transmits (hence no lines are shown between this interface and the ELANs in FIG. 6).

Switching Based On Priority

ATM physical interfaces may also be assigned relative priorities and switched on that basis. An option in router 12 may be set, e.g., by a user, to trigger automatic switching based on priority. The priority information upon which switching is based may be included in interface redundancy group information 48, for example.

Referring to FIG. 7, a process 78 is shown that is performed by interface switching code 44 to switch interfaces based on priority. Process 78 designates 80 any number of physical interfaces on a sliding priority scale, from highest priority to lowest priority. In router 12 (FIG. 2), only two physical interfaces are shown. For the sake of illustration, therefore, physical interface 26a is designated 80 as high priority and physical interface 26b is designated 80 as low priority.

Once physical interfaces 26a and 26b are designated in terms of priority, switching between them may be performed. Process 78 monitors 82 high priority physical interface 26a to determine if it is up and running. If high priority physical interface 26a is available 84 (i.e., it is not "down"), process 78 switches 86 from low priority physical interface 26b to high priority physical interface 26a. As a result, the "best available" interface on router 12 is used to provide network layer connections over network 24. Process 60 may be incorporated into processes 56 and 66 to provide further enhanced switching capabilities.

As noted above, the invention may be used on routers having more than two physical interfaces. For example, in FIG. 8, router 90 has three physical interfaces 92a, 92b and 92c organized into two interface redundancy groups 94 and 96. Interfaces 92a and 92c are in group 94 and interfaces 92b and 92c are in group 96. In this example, interface 92c

is a backup for interfaces 92a and 92b. If either of interfaces 92a and 92b fail, then interface 92c will assume the configuration of either (or both of) failed interfaces 92a and 92b.

5 Switching in the case of more than two interfaces is performed in the same way as described in processes 56, 66 and 78, except that a decision must be made regarding which (lower priority) interface to switch for the primary interface. This decision is made based on interface
10 availability and the priority of the various interfaces.

 The invention is not limited to the specific hardware and software configurations described herein. For example, the invention can be used outside the context of ATM WANs, and with network devices other than routers. In
15 this regard, it is to be understood that while the invention has been described in conjunction with the detailed description thereof, the foregoing description is intended to illustrate, and not to limit, the scope of the invention. Other aspects, advantages, and modifications are within the
20 scope of the following claims.

What is claimed is: